

CLAIMS

What is claimed is:

[0189]

[Note: Bracketed **bold and italicized** cross-referencing text is provided in the below claims as an aid for readability and for finding corresponding (but not limiting) support in the specification. The bracketed text is not intended to add any limitation whatsoever to the claims and should be deleted in all legal interpretations of the claims and should also be deleted from the final published version of the claims.]

[0190]

1. A transmission system [103, *Fig. 3*] for sending payload data [534] from a payload source [301] to a switch fabric layer [105], said transmission system comprising:
 - (a) a grant-stamp associator [335] which associates a first Grant Time Stamp (GTS-a) [325] obtained from the switch fabric layer with corresponding payload data that is to be passed through the switch fabric layer during a time slot represented by the first Grant Time Stamp; and
 - (b) a ZCell-builder [336, 310], responsive to the grant-stamp associator, and adapted to build a ZCell data structure [501] having a payload section [520] which includes a second Grant Time Stamp (GTS-b) [522], where said GTS-b is copied from, or derived from said GTS-a and where the payload

15 section further includes a copy [534] of said corresponding payload data that is logically associated with the second Grant Time Stamp (GTS-b).

[0101] 2. A transmission system [103, *Fig. 4*] for sending grants [325] responsive to requests for forwarding payload data [534] from a payload source [301] to a switch fabric layer [105], said transmission system comprising:

5 (a) a grant-stamp generator [321] which associates with a winning request [351.3x8], a first Grant Time Stamp (GTS-a) that defines a future time slot won by the winning request for moving corresponding payload data through the switch fabric layer during the time slot represented by the first 10 Grant Time Stamp; and

(b) a ZCell-builder [412], responsive to the grant-stamp generator, and adapted to build a ZCell data structure [501] having a control section [510] which includes the first Grant Time Stamp (GTS-a) [575].

[0102] 3. A manufactured signal [501] structured for transmission between a switch fabric layer [105] and a line

card layer [101], said manufactured signal comprising at least two of:

5 (a) a dual-use, request and grant field [514] for carrying one or more routing requests when moving from the line card layer to the switch fabric layer, and for carrying grant information when moving from the switch fabric layer to the line card layer;

10 (b) at least for when moving from the line card layer to the switch fabric layer, a combination of a payload-containing field [534] and another field [522] for carrying a payload-associated, Grant Time Stamp (GTS-b), where the GTS-b identifies a time slot within a destination part of the 15 switch fabric layer during which the associated payload will be switched for egress to a request-defined one or more parts of the line card layer;

20 (c) at least for when moving from the switch fabric layer to the line card layer, a combination of a source identifier (SLIN) [526] and a payload sequence identifier [525] for respectively identifying a part of the line card layer from which the payload ingressed into the switch fabric layer and for identifying a spot within a sequence of payloads at which the line card layer-carried payload is to 25 be disposed; and

(d) an error checking and correcting field (ECC) [545] adapted for use in DC-balanced transmission paths and covering included ones of items (c) and (d).

[0193] 4. The manufactured signal [501] of Claim 3 and further comprising:

(e) a back pressure field [512] adapted for use both when moving from the line card layer to the switch fabric 5 layer, and when moving from the switch fabric layer to the line card layer so that the upstream one of said layers can signal the downstream other of said layers about a congestion condition developing in the upstream layer.

[0194] 5. The manufactured signal [501] of Claim 3 and further comprising:

(e) at least for when moving from the switch fabric layer to the line card layer, an egress indicator [484-491] 5 tagged on with said a combination (c) of the source identifier (SLIN) [526] and the payload sequence identifier [525], for identifying a switch fabric egress port [E1-Em] from which at least part of the manufactured signal egressed from the switch fabric layer.

[0185] 6. The manufactured signal [501] of Claim 3 and further comprising:

(e) a first cyclical redundancy field [515] for verifying an errorless state of a respective first part [510] but not all of said manufactured signal.

[0186] 7. The manufactured signal [501] of Claim 6 and further comprising:

(f) a second cyclical redundancy field [535] for verifying an errorless state of a respective second part [520] of said manufactured signal, where the second part is mutually exclusive of said first part [510] of the manufactured signal.

[0187] 8. The manufactured signal [501] of Claim 7 wherein:

(d.1) said error checking and correcting field (ECC) [545] covers the second part [520] of the manufactured signal but not the first part [510].

[0188] 9. The manufactured signal [501] of Claim 3 and further comprising for passing from an upstream part of the

line card layer to a downstream, other part of the line card layer by passage through said switch fabric layer, at least 5 one of:

(e) a congestion indicator field [527] for reporting to the downstream part of the line card layer, a congestion condition located upstream of the downstream part of the line card layer;

10 (f) an end of packet indicator field [528] for reporting to the downstream part of the line card layer that the accompanying payload-containing field [534] contains data representing an end of a line card packet;

15 (g) a start of packet indicator field [529] for reporting to the downstream part of the line card layer that the accompanying payload-containing field [534] contains data representing a start of a line card packet; and

(h) a quality of service indicator field [532] for reporting to the downstream part of the line card layer a 20 quality of service that has been bargained for by the upstream part of the line card layer.

[0100] 10. The manufactured signal [501] of Claim 3 and further comprising for maintaining consistent identification of message flows respectively in an upstream part of the line

5 card layer and a downstream, other part of the line card
layer:

(e) a flow identification field [531] that is filled in
[493] with a respective, downstream flow identification
number of the downstream, other part of the line card layer
when reaching the downstream part, where the respective,
10 downstream flow identification number corresponds by lookup
with a respective upstream flow identification number of the
upstream part.

[0200] 11. The manufactured signal [501] of Claim 3 and
wherein when said dual-use, request and grant field carries
one or more routing requests [514B,C], said request and grant
field comprises:

5 (a.1) a multicast mode flag [550,560] for indicating
whether the carried one or more routing requests is for a
multicast transmission or not [M=0]; and

(a.2) a first validity flag [551,561] for indicating
whether the carried one or more routing requests are to be
honored in the switch fabric layer or not.

[0201] 12. The manufactured signal [501] of Claim 3 and
wherein said at least two of the: (a) dual-use, request and

grant field [514]; (b) payload-containing field [534] and GTS-b carrying field [522]; (c) source identifier (SLIN) [526] and payload sequence identifier [525]; and (d) error checking and correcting field (ECC) [545] are contained in a succession of 79 or fewer data bites, where a bite includes a 10-bit character.

[0202] 13. The manufactured signal [602] of Claim 3 and wherein said at least two of the: (a) dual-use, request and grant field [614]; (b) payload-containing field [634] and GTS-b carrying field [622]; (c) source identifier (SLIN) [626] and payload sequence identifier [625]; and (d) error checking and correcting field (ECC) [645] are contained in a succession of 69 or fewer data bites, where a bite includes a 10-bit character.

[0203] 14. A succession [313,314,318] of manufactured signals according to Claim 3 and further comprising:
(e) a succession of optical channel synchronization bites [317] interspersed between the manufactured signals.

[0204] 15. The succession of manufactured signals and of optical channel synchronization bites according to Claim 14 wherein said optical channel synchronization bites are sequentially coded and include at least the K28.5 and K28.1 fiber channel characters.

[0205] 16. The succession of manufactured signals and of optical channel synchronization bites according to Claim 14 and further comprising:

5 (f) a succession of idle bites **[417]** interspersed between the manufactured signals and the optical channel synchronization bites.

[0206] 17. The succession of signals according to Claim 16 wherein said idle bites include at least the K28.0 fiber channel character.

[0207] 18. The succession of signals according to Claim 16 wherein the number of said idle bites **[417]** is greater when being transmitted **[149b]** from the switch fabric layer **[105]** to the line card layer **[101]** than when being transmitted **[149a]** from the line card layer to the switch fabric layer.

[0205] 19. The manufactured signal [501] of Claim 3 and wherein when said dual-use, request and grant field [514B] comprises:

5 (a.1) a primary egress identifying field [553] identifying an egress path desired by a corresponding primary routing request;

10 (a.2) a primary priority code field [552] associating a competitive request priority with the primary egress identifying field;

15 (a.3) a secondary egress identifying field [557] identifying an egress path desired by a corresponding secondary routing request;

20 (a.4) a secondary priority code field [556] associating a competitive request priority with the secondary egress identifying field;

(a.5) a first validity flag [551] indicating whether the carried one or more routing requests are to be honored in the switch fabric layer or not; and

(a.6) a second validity flag [555] indicating whether the secondary routing request is to be honored in the switch fabric layer or not.

[0209] 20. The manufactured signal **[501]** of Claim 3 and wherein when said dual-use, request and grant field **[514D]** comprises:

- (a.1) a Grant Time Stamp **[575]**; and
- (a.2) an associated Grant Label **[574]**.

[0210] 21. A switching method comprising:

(a) in a switch card layer **[101]**, loading flow contents into respective ones of Virtual Output Queues (VOQs), where each VOQ **[301]** is associated with a respective unicast 5 destination or a prespecified set of multicast destinations;

(b) conducting bidding competitions **[365]** between subsets of the VOQ contents to determine which of one or more smaller number of VOQ contents will be allowed to submit a passage request **[315]** to a subset-associated part **[351.3]** of 10 a switching fabric layer; and

(c) stuffing bid-winning ones of the passage requests as primary requests into respective ZCell signals **[318]** for transmission to the subset-associated parts of the switching fabric layer.

[0211] 22. The switching method of Claim 21 and further comprising:

5 (d) first converting [312] the request-stuffed ZCell signals to a serialized optical transmission domain, adding ECC fields [545] and inserting synchronization bites [317]; and

(e) transmitting the first converted ZCell signals with ECC fields and synchronization bites by way of optical medium [103c,445] in an interface layer [103'] to the switching fabric layer.

[0212] 23. The switching method of Claim 22 and further comprising:

5 (f) second converting [123] the request-stuffed and optically-transmitted ZCell signals to a less-serialized electronic processing domain [8bpc];

10 (g) in the switch fabric, conducting grant competitions [321] between received requests associated with said VOQ contents to determine which of one or more of VOQ's will be allowed to submit a smaller number of respective payloads for passage through a grant-associated part [351.3] of the switching fabric layer and at what allocated time slots; and

(h) injecting [325'] grants and corresponding first Grant Time Stamps (GTSt) into respective ZCell signals [425]

for transmission back to the request-associated parts of the line card layer.

[0213] 24. The switching method of Claim 23 and further comprising:

5 (i) third converting **[412]** the grant-carrying ZCell signals to more-serialized optical transmission domain format **[10bpc]**, adding ECC fields and inserting sync bites and idle bites **[417]**; and

(j) transmitting **[445]** the third converted ZCell signals with ECC fields and sync bites and idle bites by way of optical medium **[445]** in the interface layer **[103']** to the switch card layer.

[0214] 25. The switching method of Claim 24 and further comprising:

5 (k) fourth converting **[483]** the grant-carrying ZCell's to the less-serialized electronic processing domain **[8bpc]**; and

(l) in the line card layer, inserting **[336]** grant-winning payloads and associated second Grant Time Stamps (GTSb) into respective ZCell signals for transmission back

[337] to the grant-giving parts of the switching fabric layer.

[0215] 26. The switching method of Claim 25 and further comprising:

(m) fifth converting [312,102] the payload-carrying ZCell signals to the optical transmission domain [10bpc],
5 adding ECC fields and inserting sync bites; and
(n) transmitting the fifth converted ZCell's with ECC fields and sync bites by way of optical medium [103a] in the interface layer to the switching fabric layer.

[0216] 27. The switching method of Claim 26 and further comprising:

(o) sixth converting [103] the payload-carrying ZCell signals to the electronic processing domain [8bpc];
5 (p.1) in the switch fabric layer, re-aligning [437,438] the ZCell-carried payloads according to their respective, second Grant Time Stamps (GTSt); and
(p.2) switching [455] the re-aligned payloads through the switch fabric layer during time slots associated with their respective, second Grant Time Stamps (GTSt).

[0217] 28. The switching method of Claim 27 and further comprising:

5 (q) seventh converting **[1R2]** the switched payload-carrying ZCell signals to the optical transmission domain **[10bpc]**, adding ECC fields and inserting sync bites and idle bites; and

(r) transmitting **[419b]** the seventh converted ZCell signals with ECC fields and sync bites and idle bites by way of optical medium **[445]** in the interface layer to the line card layer.

[0218] 29. The switching method of Claim 28 and further comprising:

5 (s) eighth converting **[1J3,483]** the switched-payload-carrying ZCell signals to the electronic processing domain **[8bpc]**; and

(t) in the line card layer, re-ordering **[487]** received ones of the switched-payloads according to accompanying source and sequence designations.

[0219] 30. The switching method of Claim 29 and further comprising:

(u) attaching [493] destination-based flow identification numbers (FIN) to the re-ordered and switched-
5 payloads; and

(v) forwarding [498] the FIN-bearing switched-payloads to their respective destination lines.

[0220] 31. A switch fabric integrated circuit [151, ~~zest~~] comprising:

(a) a request receiver [315] adapted to receive requests for access to switching resources [255, ~~455~~] within the 5 integrated circuit, from plural requesters, the received requests each including a respective priority code [552] and an identification [553] of the requested switching resources;

(b) a request queue [411] for storing received requests and allowing a tracking of aging of the stored requests;

10 (c) a request arbitrating mechanism [321] for arbitrating amongst received requests that seek to access same switching resources [329], said request arbitrating mechanism being responsive to:

(c.1) busy states [372] of requested resources;

15 (c.2) request aging; and

(c.3) request priority code; and

(d) a grant returning mechanism [325], operatively coupled to the request arbitrating mechanism, for sending grant signals [514D] back to respective requesters whose 20 requests won grants in the request arbitrating mechanism, the grant signals each having a first Grant Time Stamp (GTS-a) [575] identifying a future time slot [439] of the switch fabric integrated circuit in which the requested resources will be made available, and a grant label [574] identifying the granted resource [voq#].

[0221] 32. A switch fabric interfacing, integrated circuit (IC) [119,ZINC] comprising:

5 (a) a payload output queue [301] for storing received payloads;

(b) a bid distributing and arbitrating mechanism [310] for distributing payload bids to localized arbitration pools [365] and for arbitrating amongst the distributed bids to decide which of the bids will send requests [315] that seek to access of switching resources [455,329] of a switch fabric, 10 said bid distributing and arbitrating mechanism being responsive to:

(b.1) payload priority codes [552]; and

15 (b.2) depth of output queue fill for competing
ones of the payloads stored in the payload output
queue; and

(c) a request inserter [312] for inserting requests
[315] of winning bids into request-transporting cells [318]
moving out of the switch fabric interfacing IC.

[0222] 33. A protocol conversion method comprising:

(a) receiving in a source line card [110], payload data
that is transmitted according to a first transmission
protocol [115];

5 (b) re-arranging the received payload data for transport
in payload-carrying sections [314,534] of intermediate
transmission signals (ZCells) [501];

(c) transmitting the re-arranged payload data along with
dynamically-assigned, Grant Time Stamps (GTSb's) [522] to a
10 switching chip [151];

(d) in a time slot designated by the carried-along Grant
Time Stamp (GTSb), switching [445] the re-arranged payload
data through the switching chip;

(e) transmitting the switched payload data along with
15 associated source and sequence designations [526,525] to a
line card chip [139] of a destination line card [130]; and

(f) in the destination line card, re-arranging the switched and transmitted payload data for further transmission according to a second transmission protocol [145] that is different from the first transmission protocol.

[0223] 34. An interface layer [103', ~~Fig1c~~] for provision between a line card layer [101] and a switch fabric layer [105] comprising:

5 (a) first converting means [152] for receiving payload-carrying ZCell signals coded in parallel in an electrical transmission domain [8bpc] and for converting the received ZCell signals into a serialized optical or electrical transmission domain [10bpc], for adding ECC fields [545] to the converted ZCell signals, and for inserting sync bites 10 between the ones or plural sets of the converted ZCell signals;

15 (b) serializing means [155], coupled to the first converting means, for serializing the converted ZCell signals and the inserted sync bites;

(c) serial transmission means [103c], coupled to the serializing means, for optically and/or electrically transmitting the serialized ZCell signals and inserted sync bites;

20 (d) deserializing means [**1R5**], coupled to the serial transmission means, for deserializing the optically and/or electrically transmitted ZCell signals and inserted sync bites; and

25 (e) second converting means [**1R3**] for receiving the deserialized and payload-carrying ZCell signals, for using the ECC fields [**545**] to test the received and deserialized ZCell signals for errors and correct errors, for using the deserialized and inserted sync bites to synchronize the received and error checked/corrected ZCell signals with a local clock, and for converting the received ZCell signals into the electrical transmission domain [**8bpc**].

[0224] 35. The interface layer [**103', Fig 1C**] of Claim 34 wherein:

(c.1) said serial transmission means [**103c**] extends between spaced apart shelves [**70,80**] each containing a respective one or more of line cards and switch cards.